This Elephant can Dan

Robert Calderbank
VP, *Information Sciences Research*
AT&T Labs, Florham Park, NJ

This is a talk on the role of information in service and network operations at AT&T.

Scale Changes the Face of Society

Introduction of the ass
line permitted Henry F
cut car prices and crea
mass market.

He also raised wages.

21st Century Communications Services

The AT&T Strategy: Integration of services across net
to high value customers

- **Customers**
- **Services**
- **Networks**

**These are the 3 essential eleme
communications services bu**

1. **Detailed knowledge of the activities and inferred needs
potential customers, and AT&T relationship with them.**

2. **The ability to operate, administer, maintain and provisi
services at premier levels.**

3. The ability to carry communications signals in a variety of wa
distances.

4. **Exact real-time knowledge of the detailed interactions a**

ta Mining and Data Warehousing in Telecommunicati

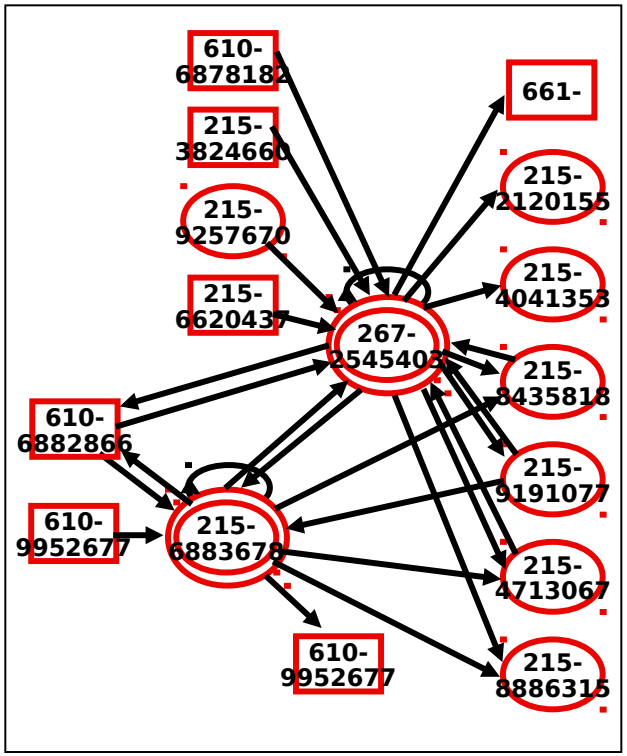

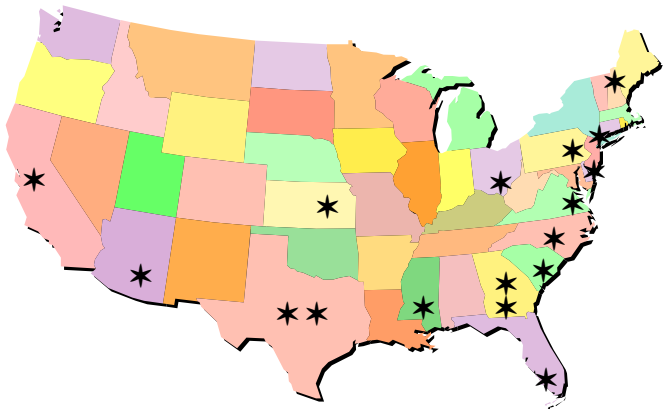

Research ⟷ **Systems** ⟷ **Operations**

What Information is Necessary to Operate and Compete
How is this Information Obtained and Applied

NETWORK

How can data be efficiently transported across Wide Area Networks?

SYSTEMS AND SOFTWARE

What types of storage and processing architectures will satisfy user needs?

DATA ANALYSIS

What information can be mined from the data?
What type of decisions can be supported?

VISUALIZATION

What are the most effective ways to deliver information to decision makers?

Application Specific Knowledge

Fraud, Customer Focused Operations
AT&T Switched Network, Frame Relay,

Data

Scale, Single Database
Scale, Multiple Databases
Scale, Multiple Databases
Multiple modes

State of the AT&T InfoLab

- **Started in 1995**
- **Approximately 50 Research**
- **Large Display Wall Installed in 1996**
- **Server Cycles (2001): > 90 GIPS, 40 GFLOPS**
- **On-Line Data (2001): 60 => 100 Terabytes**
- **Real Time Data (2000): > 50 Gigabytes per day**

Data Sources

- **Call Detail (All Distance, All Services)**
- **Frame Relay / ATM (All Data)**
- **IP Backbone**
- **IP Services**
- **Customer Care Call Centers**
- **Order / Provisioning / Billing (Order to Cash Process)**
- **And many others**

APPLICATIONS

CUSTOMERS

BUSINESS MANAGEMENT

SERVICES/ OPERATIONS & MANAGEMENT

NETWORK OPERATIONS & MANAGEMENT

NETWORK

NETWORK ELEMENTS & ELEMENT MANAGERS

CUSTOMER CARE

BILLING

SERVICE ORDERING & PROVISIONING/ OPERATIONS/ MAINTANENCE

NETWORK OSS PROVISIONING/ OPERATIONS/ ADMINISTRATION/ MAINTANENCE

ELEMENT MANAGEMENT (E.G. SNMP, CMIP, ALARM GENERATION)

MARKETING & SALES

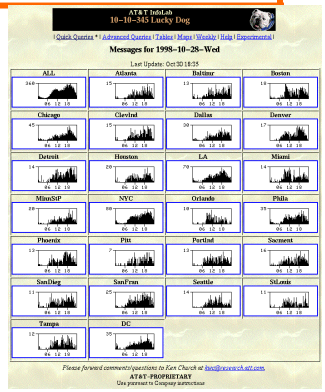

NETWORK PLANNING & ENGINEERING

USAGE RECORDING

BAD DEBT

UNCOLLECTABLES

FRAUD

Data Integrity + Business Control

800 Instant Marketing

"Technology is not the carriers battlefield. The real issue is the back office. This is where the battle will be won or lost over the next several years."

John Sidgemore

Chief Operating Officer, WorldCom

Current Real Time Service and Customer Management Architecture

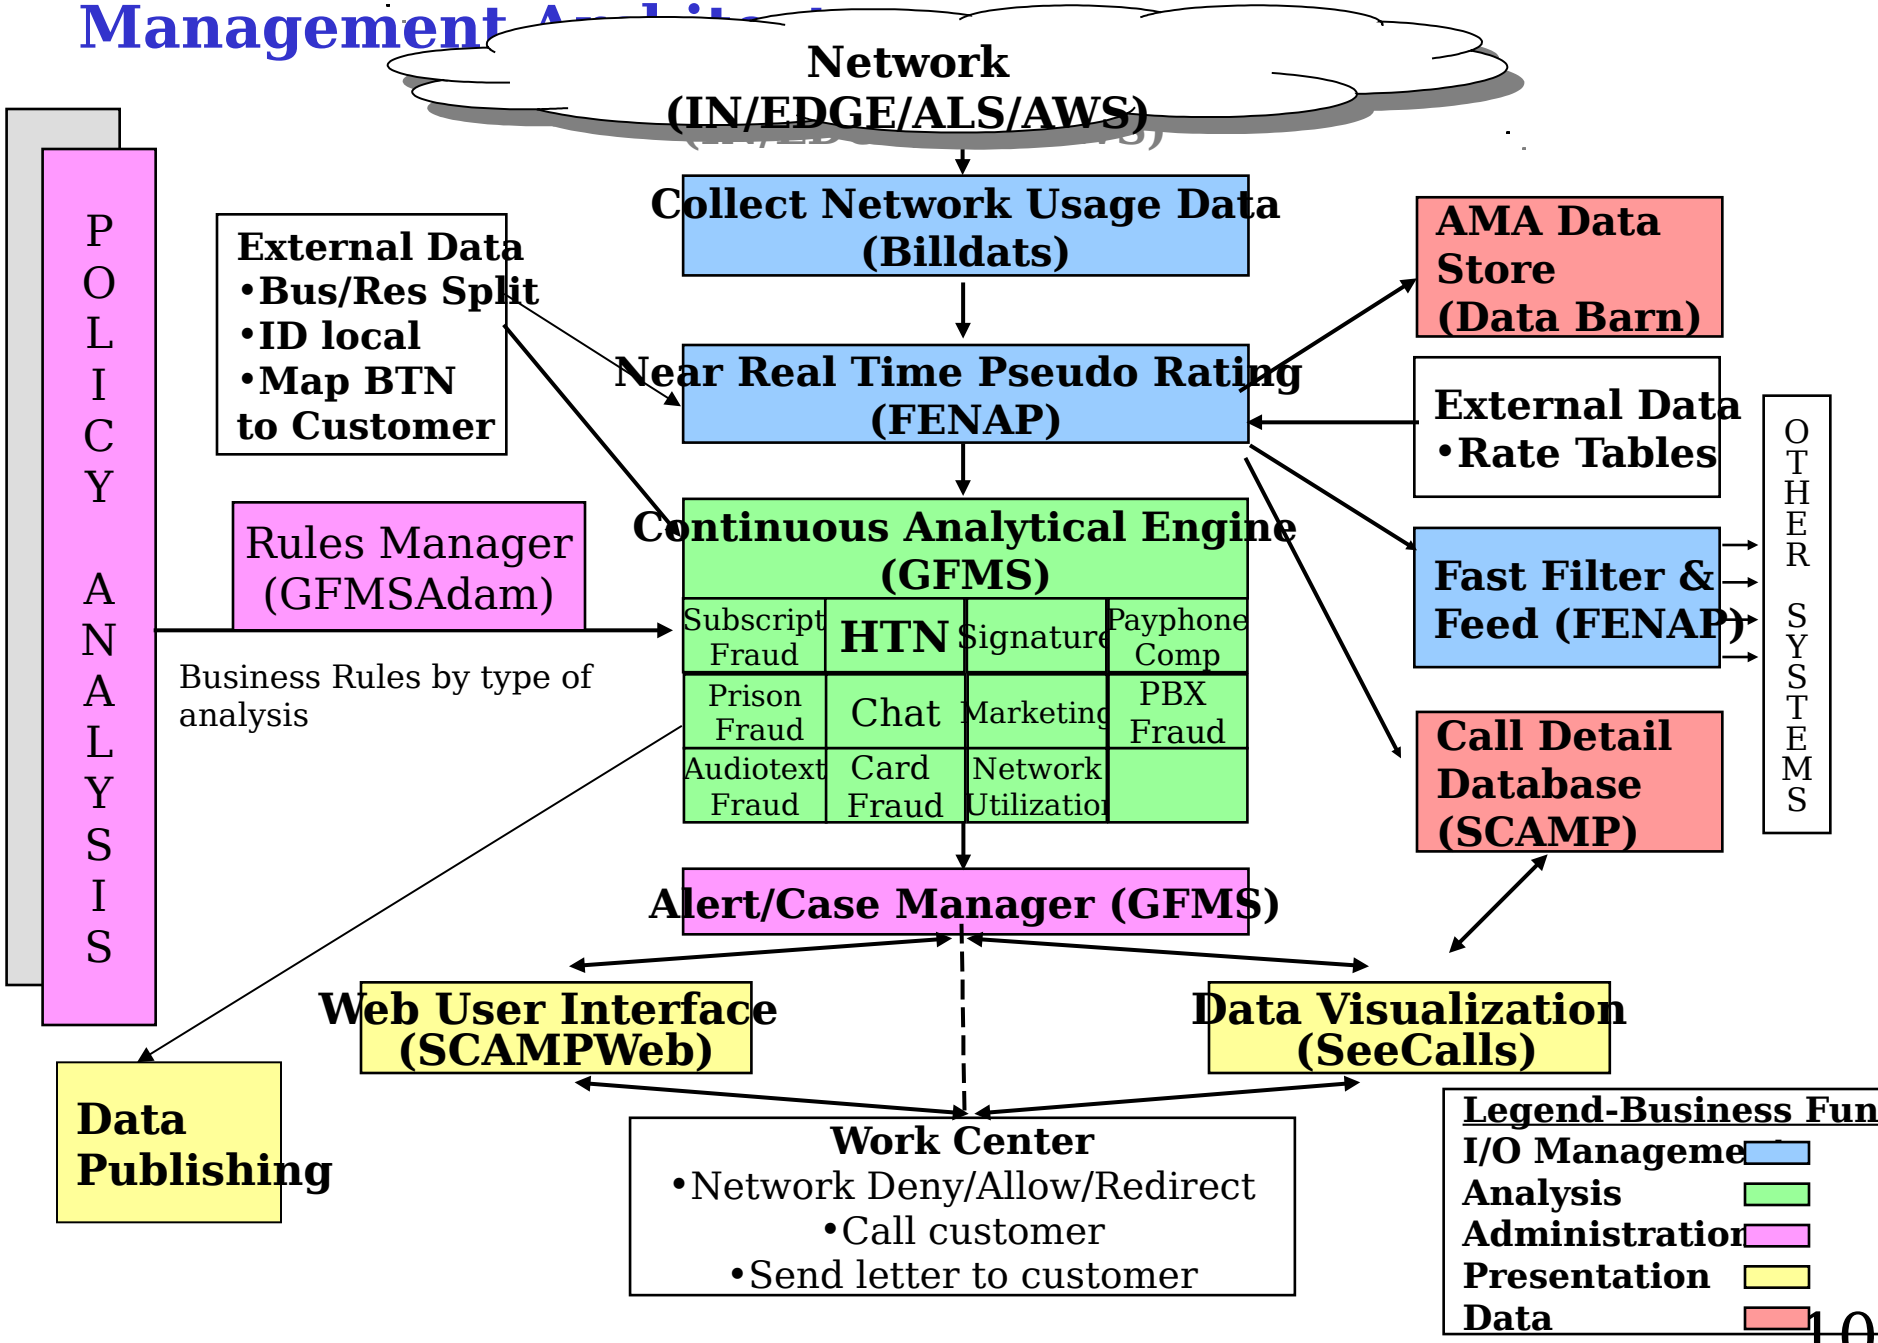

nt Real Time Service and Customer Management To

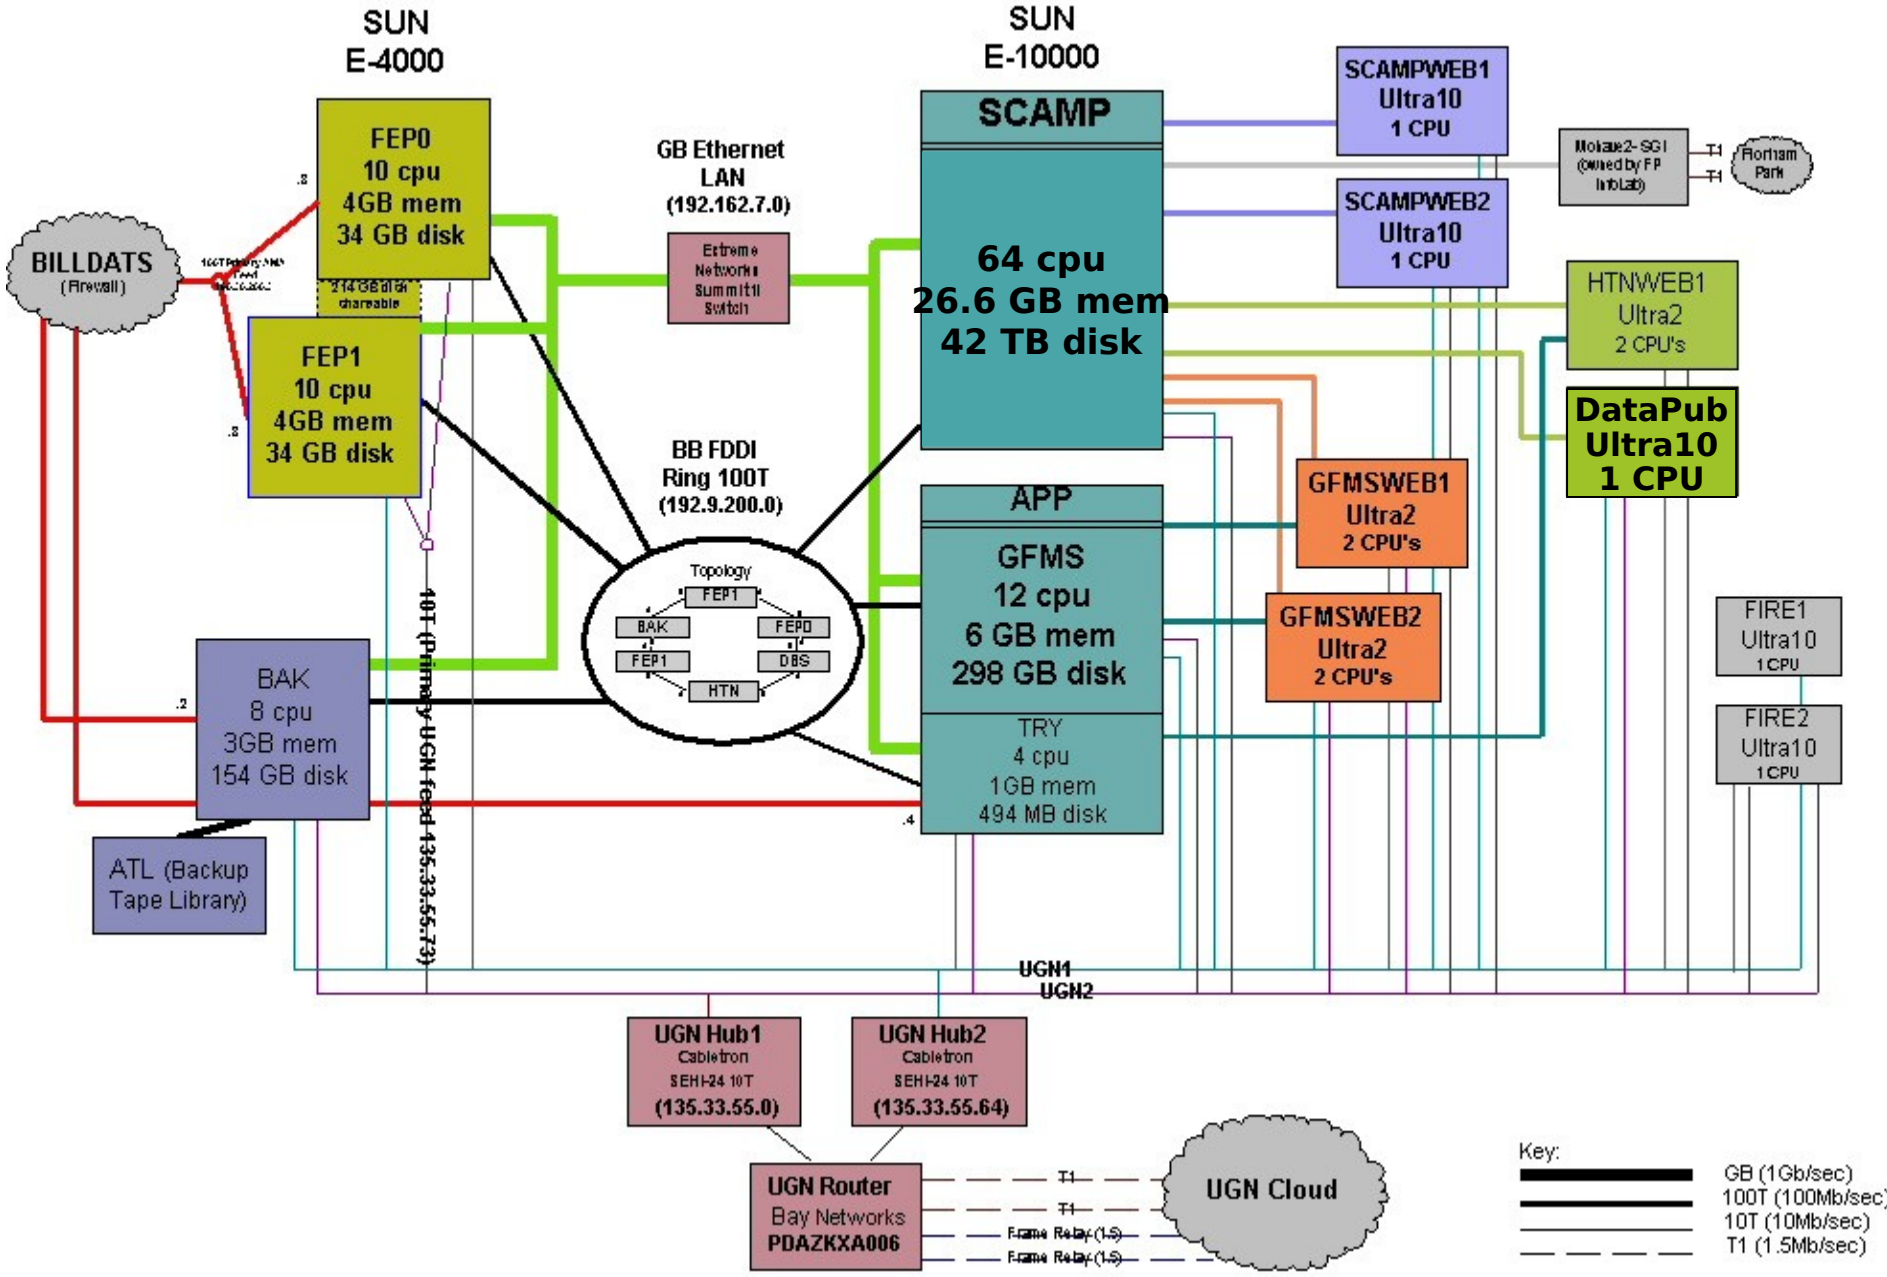

nt Real Time Service and Customer Management To

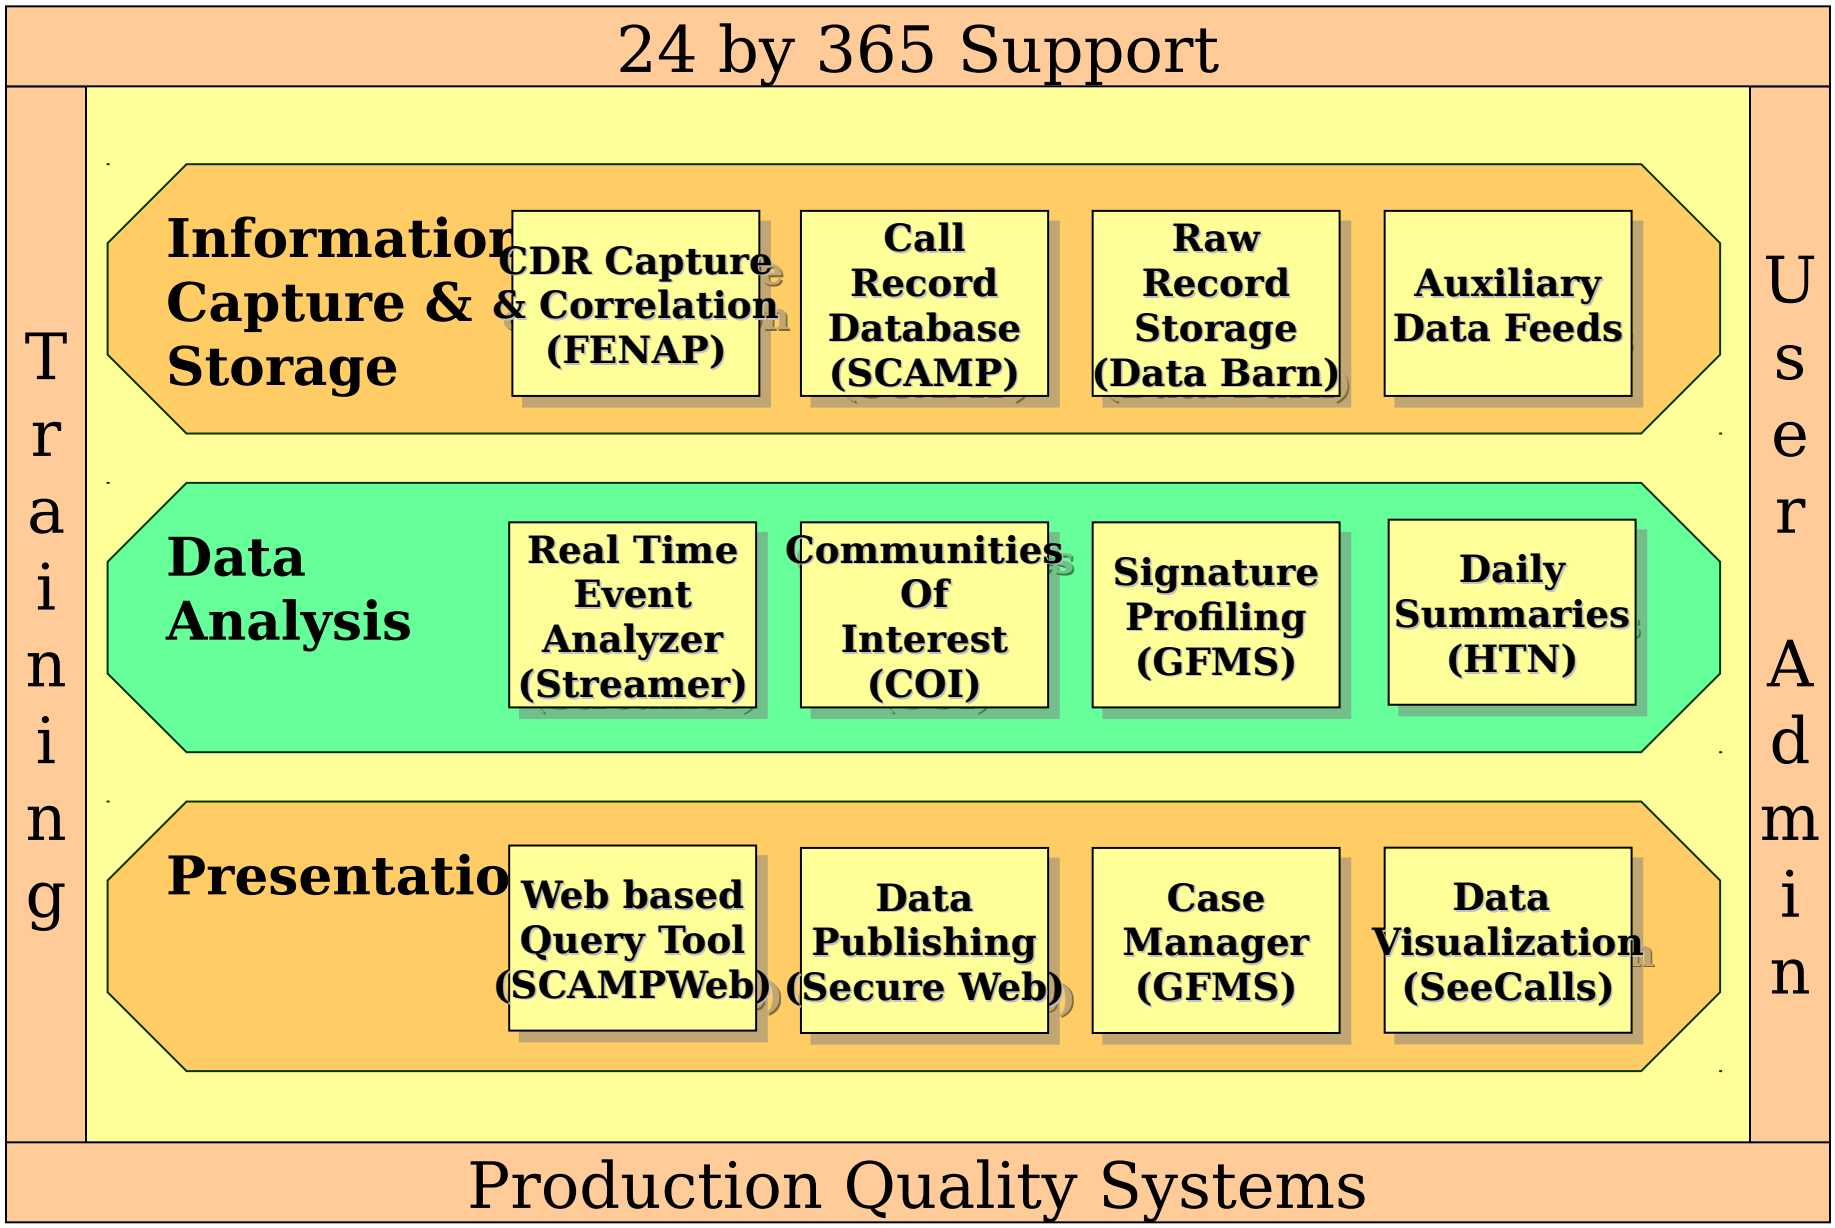

Technical Capabilities

Subscriber Base: 1500 users in 24 geographically separate locations

Database: 75 Terabytes of all AT&T voice calls (local and LD) made since 7/00

Query Response Time

All Calls from 212 to China – 90 seconds

All Calls from NPA-NXX-YYYY – 30 seconds

Operating Statistics

- 25 by 365 Availability
- 99.8% Up Time
- 100% Data Integrity
- Security at user level for all actions

Data Streams

Barbara Peda Business Markets
Wholesale Voice

AT&T Labs Streamer was the only system capable of providing a large carrier with a call detail feed meeting the contract specification of 25 minute latency

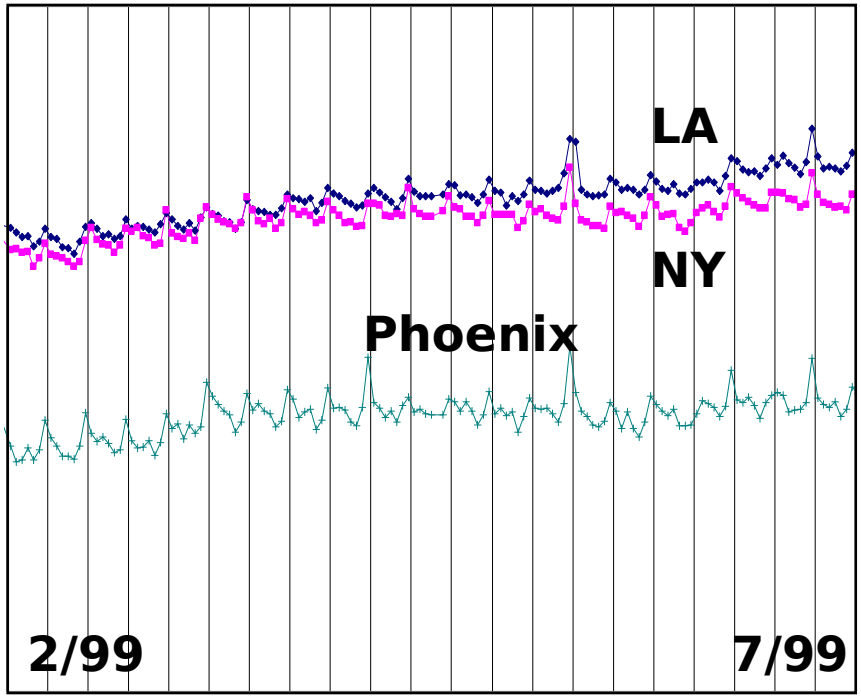

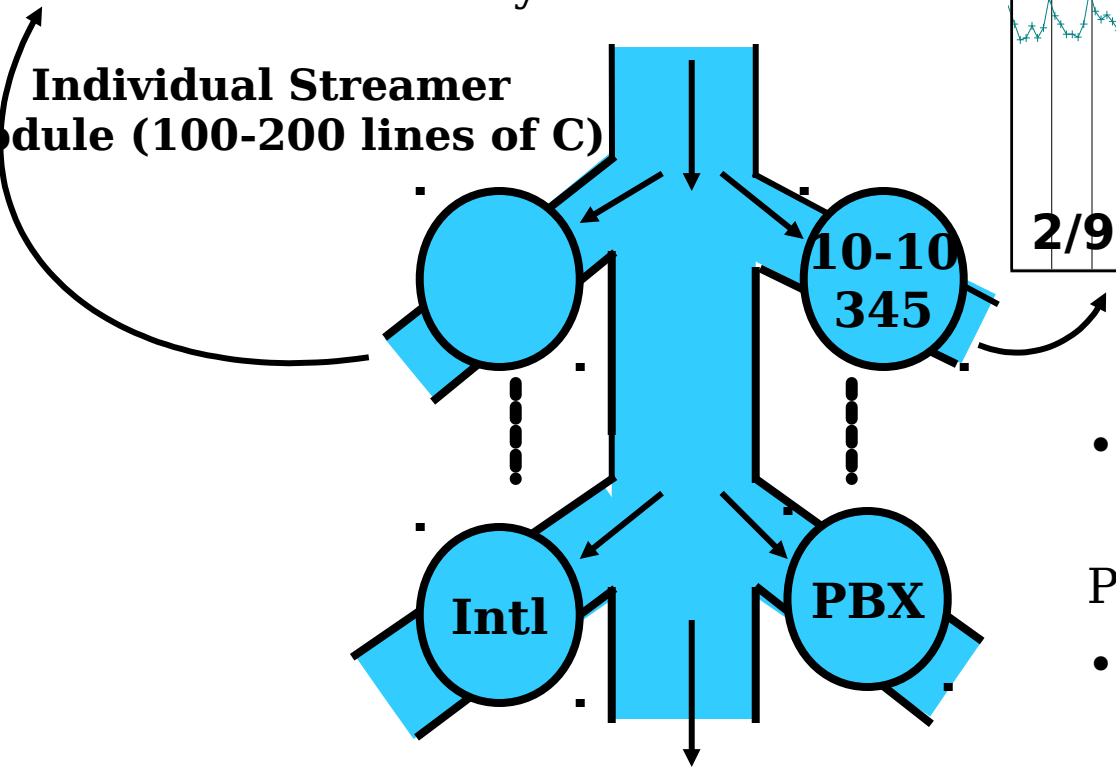

Tracking for Transaction Ser

- by geographic location
  — Los Angeles, New York, Phoenix
- by time of day or week

THE STREAMER: 200 lines of C

Understanding of the Communications Mar

TFEDS collected call detail in a window following a

GFMS reveals patterns in customer behavior

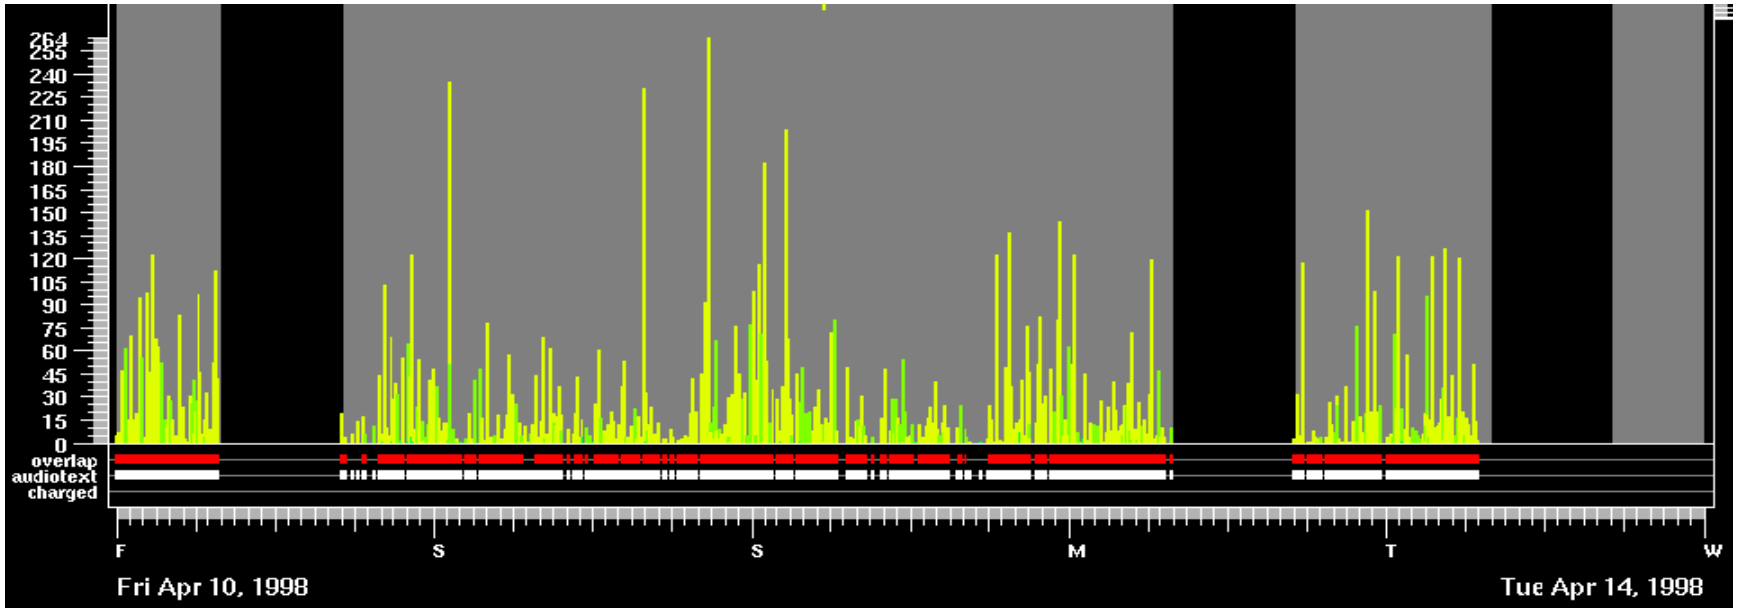

All traffic to Adult Chatlines on a Wholesale Voice Ac

CDW Re-Architecture Project Overview

Solution

- ⇒ Develop an AMA data store, Data Barn, that will store two years of AMA using pzip compression technology
- ⇒ Expand the current Daytona call detail record (CDR) data store, SCAMP, to accommodate two years worth of CDRs
- ⇒ Provide a user interface, SCAMPweb, which will look into the indexed database, SCAMP, and return the appropriate CDR as well as a pointer to the AMA

Services Provided

- ⇒ End to end monitoring, maintenance and system administration
- ⇒ Administration of all user Ids
- ⇒ Documentation and training for all systems and users

Major Milestones

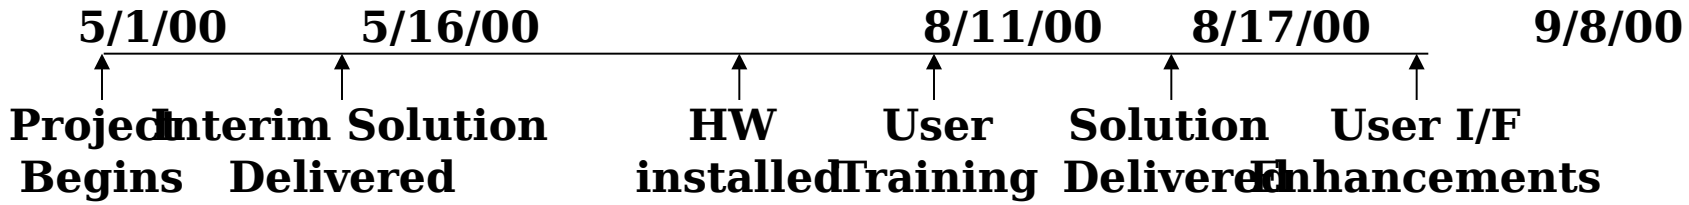

W Re-Architecture Cost/Benefit Comparis

Legacy Architecture	Data Barn Solution
⇒ $6.8M OE per year	⇒ $1M OE per year
⇒ Query response time measured in hours/days	⇒ Query response time measured in minutes
⇒ Proprietary software requires vendor support	⇒ Use of pzip compression requires less hardware
⇒ Screen based interface	⇒ Web based interface
⇒ Client server architecture	⇒ Centralized architecture

Learning From Data—AT&T InfoLab Business Operations and Network Management

- ***Data Hoarding*—measure the right things about the way the company does business, severely restrict access to the data, then store it away safely (sometimes forever), look at it only when things go wrong**
- ***Data Publishing*—make all the data within the company broadly available throughout the corporation, encourage people at all levels to study and use the data to provide competitive advantage, learn how to run the business more efficiently based on lessons learned from the data**
  - ***everyone in the organization is on the same 'web' page***

- **Understand the market first**
- **Act on changing market conditions with more insight than your competition**
- **Do it consistently**

WIN!

AT&T Network Connections Before

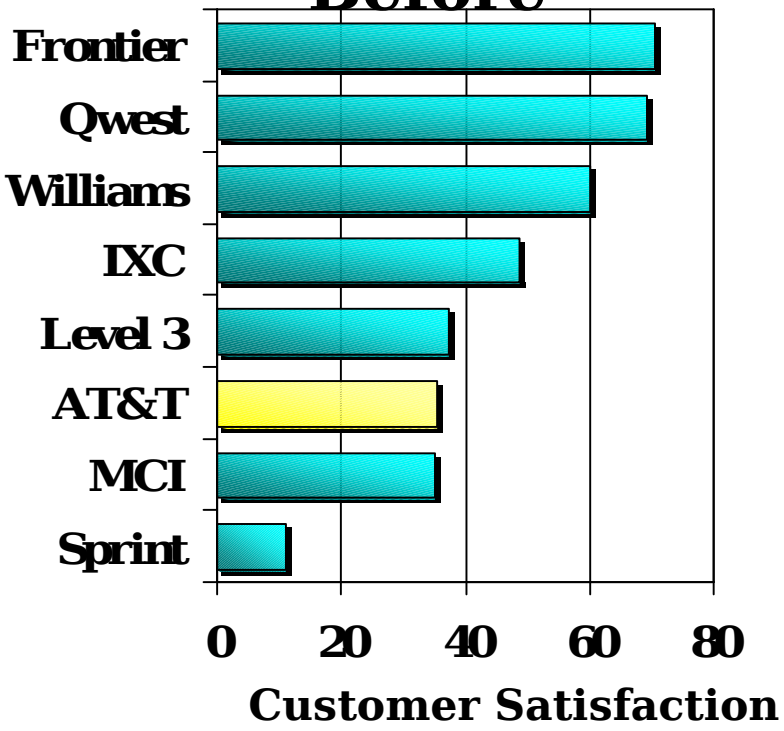

AT&T Network Connections After

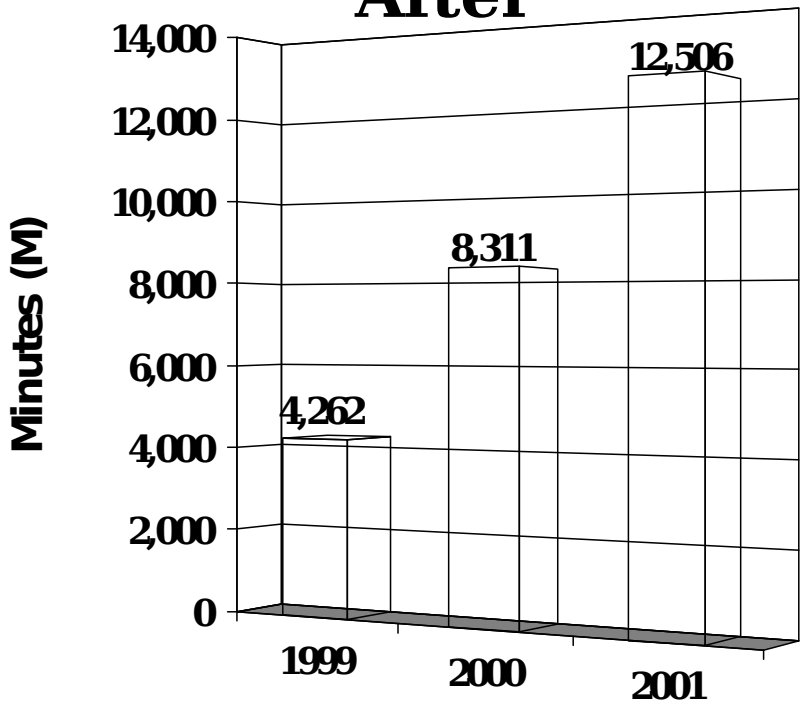

Competition Based on Price

$250M Incremental Reven

Recognition

Our new AT&T Network Connection (ANC) offer propelled us from a position of weakness in the marketplace to a position of strength. The ANC offer remains unique in the marketplace and provides a differentiator that our competitors cannot match without substantial development costs.

Peggy Sexton

Director, Marketing and Channel Development

ABS Service Providers Markets

Examples of Services

- An 800 customer runs a newspaper add in Minneapolis on April 12th. The customer is given an hourly report of all calls (minutes and messages) from Minneapolis by NPA/NXX to the 800 number featured in the ad.
- An 800 customer wants to know all the WTNs who attempted to call the 800 number but were unable to complete the call.

800 Instant Marketing Minutes Capture Market Insight

USAToday AD 2 for 1 CRUISE SALE Caribbean
Alaska Europe Panama
800-445-4177
www.great-vacations

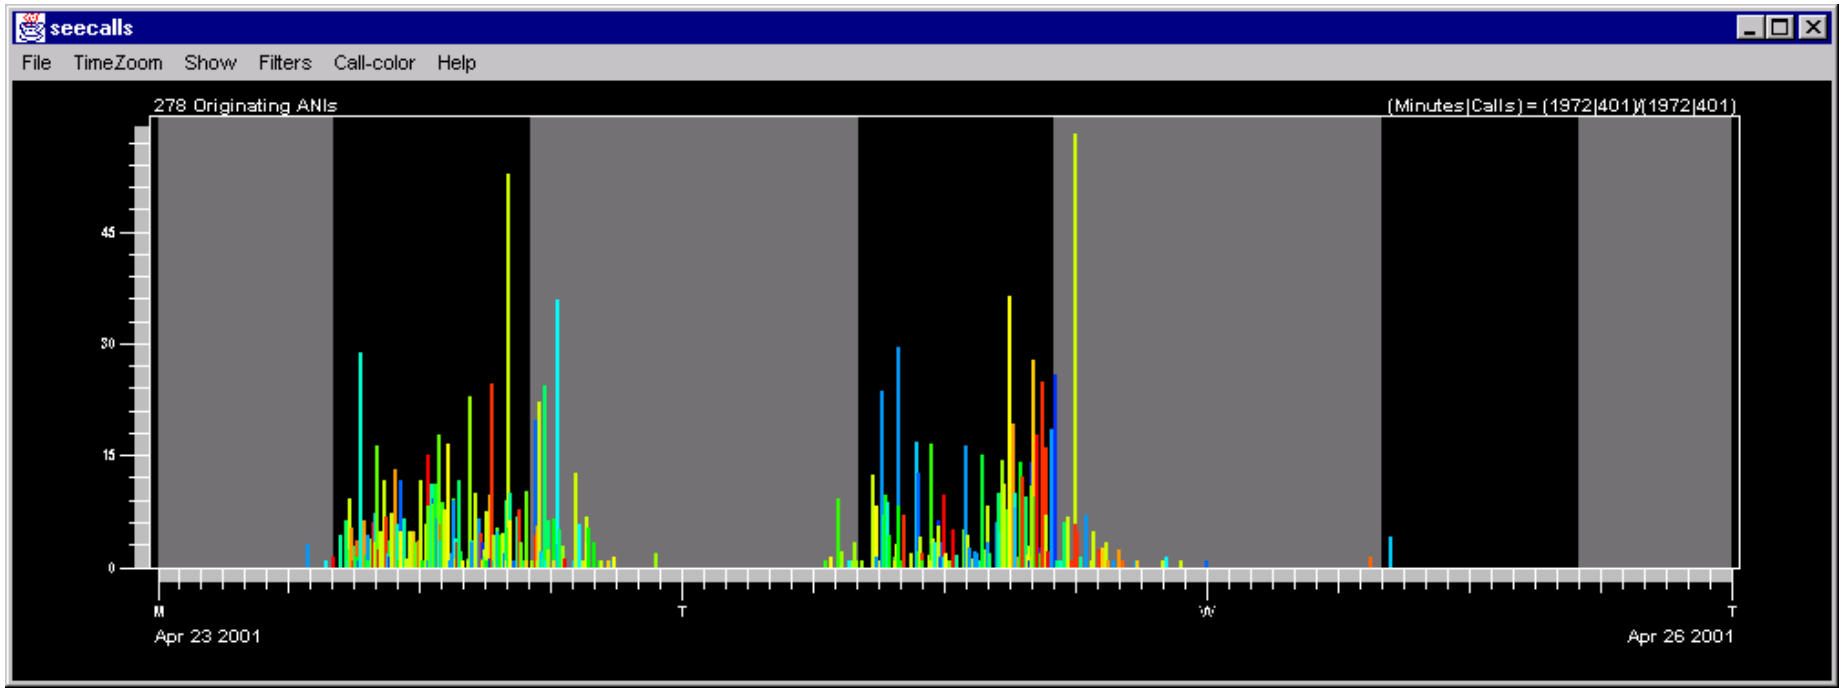

Applications for Wireless CDR's

- Fraud Detection
- Marketing
- Operational Support (e.g., identifying misrouted calls)
- Providing trunking capacity numbers to help L.A. Wireless in ordering new trunks, rings, positioning new MSCs (mobile switching centers), etc.
- Supporting the introduction of a Hub architecture into L.A. as a prototype for other AWS systems.

Exploiting Wireless Call Detail Records

- Identifying and collecting all the call detail records involved in a wireless call is nontrivial.

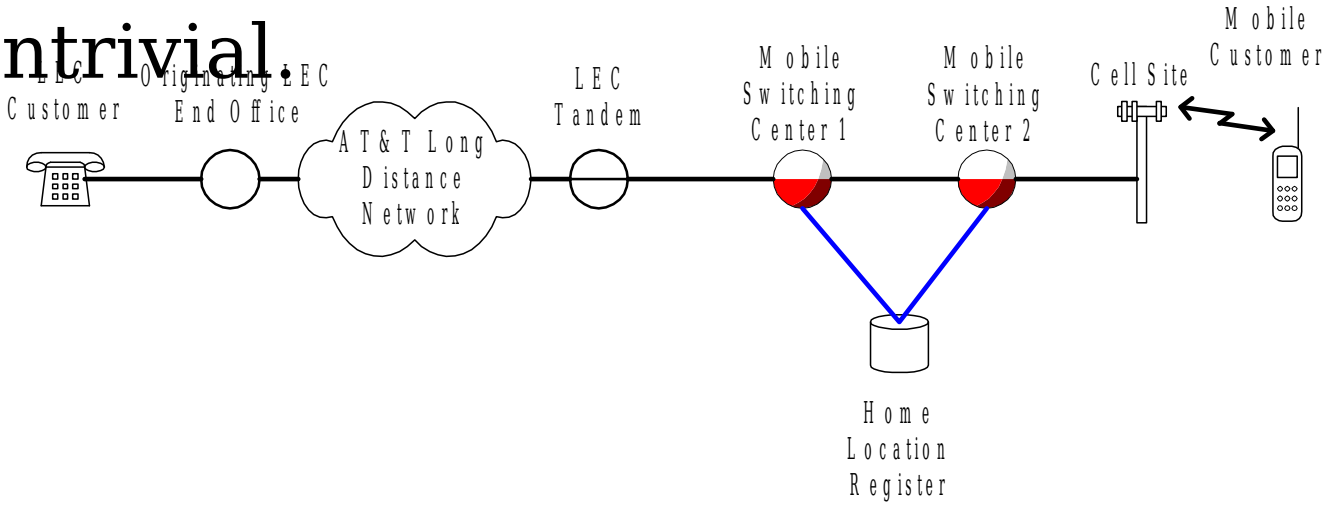

- We have developed systems for deriving and storing wireless CDR's, as well as tools for exploiting them.

Example: Raw CDR (Ericsson) Format

Record Type 2 hex	5
Call Status 1 hex	0
Record Sequence Number 8 dec	14754211
Call Identification Number 8 dec	11370051
Related Call Number 8 dec	
Cause for Output 1 dec	4
A-Number Indicator 2 dec	0
A-Subscriber Number,(LA) 28 dec*	9733607212
A-Subscriber Mobile Station Number,(LA) 10 dec	
A-Subscriber Serial Number 8 hex	
B-Number Indicator 2 dec	5
B-Subscriber Number, (LA) 28 alpha	2136109000
B-Subscriber Mobile Station Number,(LA) 10 dec	2136109000
B-Subscriber Serial Number 8 hex	7A2464B7
A-Category 2 dec	1
B-Category 2 hex	0
Date for Start of Charging 6 dec	20001018
Time for Start of Charging (24 hour clock) 6 dec	113222
Chargeable Duration 6 dec	26
Time from A-party Voice Channel Seizure to Start of Charging Start 6 dec	
Type of last voice channel of the TT record for A-subscriber 2 dec	
Internal representation of the last voice channel of the TT record for A-subscriber 5 dec	
Time from B-party Voice Channel Seizure to Start of Charging 6 dec	
Type of last voice channel of the TT record for B-subscriber 2 dec	
Internal representation of the last voice channel of the TT record for B-subscriber 5 dec	
Time from Register Seizure to Start of Charging 6 dec	14
Time of initialization of the handoff data section for the A-subscriber 6 dec	
Time of output of the handoff data section for the A-subscriber 6 dec	
Global route of the handoff data section for the A-subscriber,(LA) 7 alpha	
Time of initialization of the handoff data section for the B-subscriber 6 dec	
Time of output of the handoff data section for the B-subscriber 6 dec	
Global route of the handoff data section for the B-subscriber,(LA) 7 alpha	
Interruption Time 6 dec	0
Abnormal call release 1 dec	0
Fault Code 5 dec	
Internal Cause and Location 4 hex	203
Rerouting Indicator 1 dec	0
Exchange Identity, (LA) 15 alpha	LAXT1*311000*16
Record Number (for partial output) 2 dec	1
Outgoing Route, (LA) 7 alpha	LAXS1IT
Incoming Route, (LA) 7 alpha	LA1DORI
Tariff Class 3 dec	2
Origin for Charging 4 dec	0
Type of Procedure 2 dec	0
Subscriber Service Indicator 2 dec	0
X-Subscriber Number,(LA) 28 dec*	
Calling Line Identification Confirmation 2 hex	
Voice Mail Retrieval Indicator 2 hex	
Message Waiting Indicator 2 hex	
Originating Call Access to IN Services 2 hex	
Authentication Deny Access Indicator 2 hex	
Charging Area Information 2 hex	
ICI Indicator 1 hex	0
Output for Subscriber 1 hex	2
Roamer Port Indicator 1 hex	
Carrier Identification Code, (LA) 6 hex	
Billing ID 14 hex	001B01DC533300
Related Billing ID 14 hex	
Bandwidth 1 dec	
Service Code 1 dec	
Location Routing Number (LRN) 10 dec	

Call Type & Status

"A" Number

"B" Number

Date & Time

Fault Code

Switch Identifier

Outgoing and Incoming Routes

Billing ID Numbers

Example 1: Best Cost Routing Introduction

- Best Cost Routing involves choosing among several possible network configurations
- Examples in a wireless telephone network:
  - 2A vs. 2B routing for outbound local traffic
  - Using an IXC vs. establishing a POP to deliver traffic in another LATA (e.g., Reno backhaul)
  - Using MSCs as tandems vs. establishing dedicated routes (e.g., direct connect to neighboring market for call delivery within a WSCA)

2A/2B Example

- Can send M-L local traffic to the LEC tandem (2A) or to the LEC end office (2B)
- Costs for terminating calls via 2B are lower than for 2A, depending on interconnection agreements (60% savings for L.A. contract)
- 2B facility costs are higher because requires new dedicated trunks between MSC and end office
- Bottom line
  - Net savings of $50k per year from using 2B routing to a single busy end office
  - Possibly ~ $1M per year for all offices in a large market
- This type of Least Cost Routing allows AWS to remain profitable while being price competitive in the marketplace

2A/2B Analysis

- Need to find the busy end offices and determine optimum number of 2B trunks
- Data required:
  - Rates for 2A & 2B call termination (contract)
  - Rates for dedicated facilities, trunk charges, etc.
  - Network topology information
  - Minutes of Use (MOU) to each End Office
- Procedure for AWS-LA network design:
  - Select an End Office to study
  - Find distance to closest SONET ring hub
  - Calculate MOU to EO from Call Detail Rec's
  - Vary # of 2B trunks and calculate direct and overflow traffic
  - Sum all cost components and generate a curve
  - Find minima of curve

2A/2B Cost Components

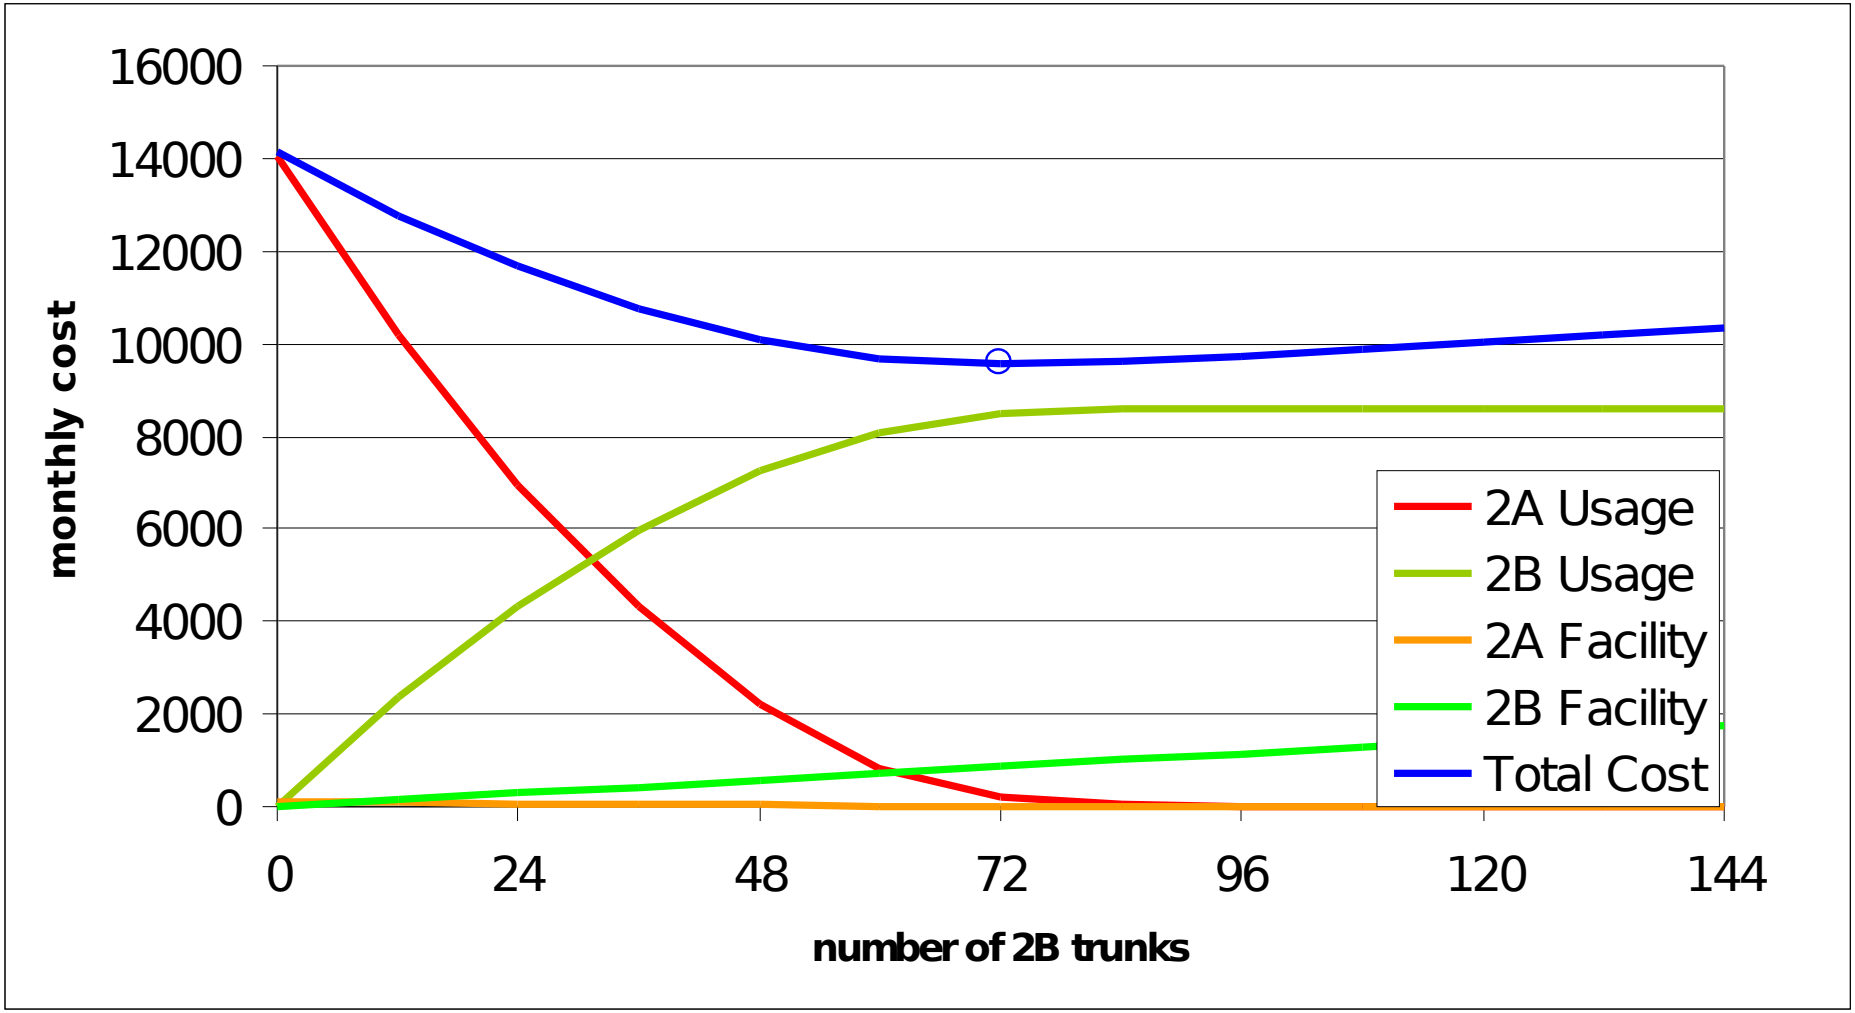

L.A. Wireless Network Re-Design

- Converting to a Gateway Architecture
- Nortel doing the design, based on our data

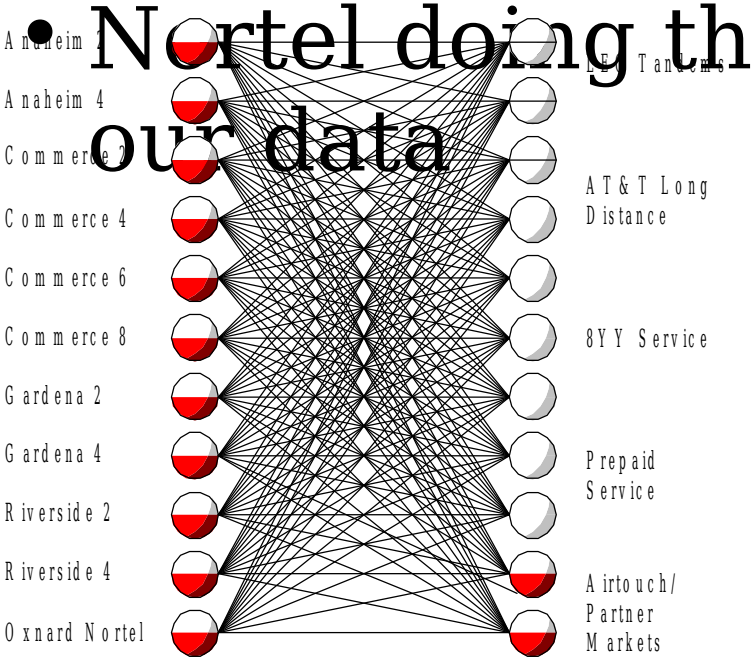

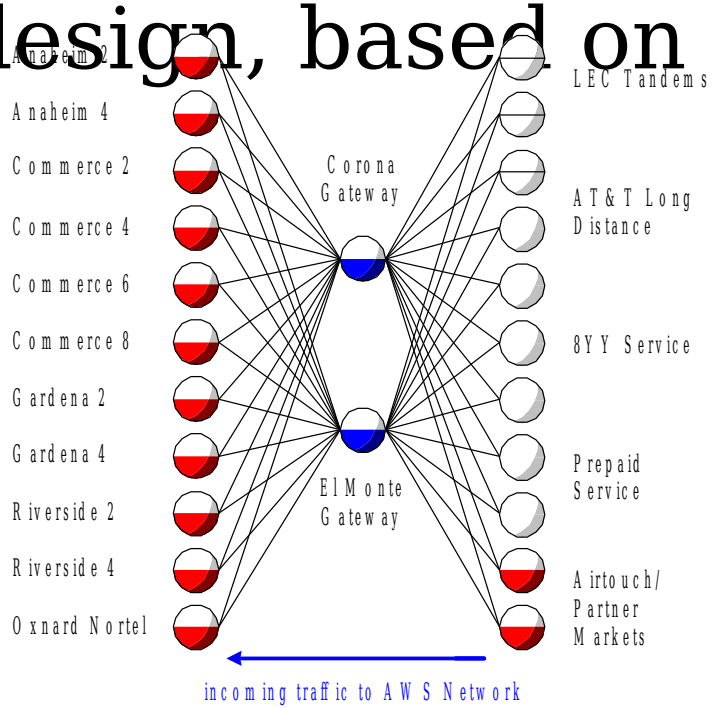

Before: Each mobile switching center (MSC) connected directly to all PSTN's

After: Connections made through two high-capacity Hubs (saves $50M over 5 years)

Questions that CDR's help answer

- In calls from PSTN to mobile, which NPA-Nxx's should go to which hubs (load balancing)?
- Given the above assignments, how much capacity should be assigned to each PSTN-Hub and Hub-MSC link?

CDR's tell us that we shouldn't simply make all trunks the same size

Nortel using our data to choose appropriate sizes

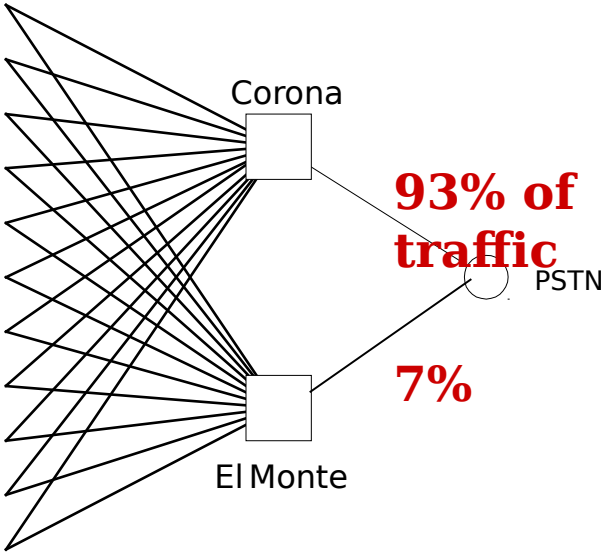

Summary

- Interaction with network experts to understand the interplay between network issues and traffic patterns
- Accurate and timely data
- Micro level detail (Switch counts not sufficient)
- Ability to handle massive amounts of data
- Dynamic models at a detailed level
- Easy to use (Point & Click capability)
- Ability to drill down
- Accessibility to raw data for validation

AWS/ AT&T Labs SMEs

SCAMP

Streamer

SCAMPWeb SeeCalls

Data Barn

BACKUP SLIDES

Professional Service Offerings

Contact:

Patricia E. Wirth

AT&T Labs

Middletown, New Jersey

732.420.2100

pwirth@att.com

Network Design & Restoration Planning

Description: A team of highly experienced telecommunications engineers will analyze an IP, ATM, Frame Relay, or Optical Transport network. They will determine where, when, and what type of facilities are needed to handle traffic demand growth and to provide resiliency to component failures. In providing this service, we use a suite of sophisticated tools, developed and used for the largest and most technologically advanced networks at AT&T. Deliverable is a network plan containing detailed short and long-term recommendations on topology, dimensioning, and restoration capacity.

Profile of Prospective Customers: Our clients are typically the network planning organization of telecommunications carriers as well as planning divisions of large enterprises who manage their own networks.

Value to Customer: The benefit is a reduction of the capital outlay required while maintaining or exceeding a desired service level.

Engagement Duration/Price: Typical projects take 3 to 6 months, costing $150K - $500K.

Contact(s): Patricia E. Wirth (732-420-2100) will provide assistance from the initial client discussion through proposal preparation and project delivery.

Service Reliability & Defects Per Million

Description: A team of highly experienced telecommunications engineers provide reliability/availability/DPM analysis for products, networks and services. Our experts work with the customer to ***define*** the reliability objectives, cost effectively ***design*** to meet the reliability objectives, implement a methodological program to ***measure*** and track the reliability objectives, ***analyze*** the metrics via Root Cause Analysis, and make recommendations to cost effectively improve reliability.

Profile of Prospective Customers: Our clients include large core IP network operators and telecom service providers, equipment manufactures, and enterprise network managers.

Value to Customer: Customer benefits are the ability to define and measure critical services and a plan of continuous improvement toward attaining cost effective reliability objectives.

Engagement Duration/Price: Our engagement stages range from initial design to analysis of field data from deployed products/services. Typical engagements last 3 to 6 months, costing $150K - $500K.

Contact(s): Patricia E. Wirth (732.420.2100) will provide assistance from the initial client discussion through proposal preparation and project delivery.

Load Characterization & Capacity Planning

Description: An expert engineer identifies key load and server usage data that need to be collected, analyzes the collected data, and uses it to develop a capacity planning model that projects the server resources required over time. Deliverables include a final report containing a summary of the data analysis and detailed short and long-term recommendations, as well as optional set of Spreadsheet-type tools.

Profile of Prospective Customers: Typical buyers are engineering and/or operations managers who proactively seek to stay ahead of the service demand in order to provide ongoing high quality service.

Value to Customer: The benefit is the ability to stay ahead of the demand growth and make new capital purchases before performance deteriorates due to excessive load.

Engagement Duration/Price: Typical projects take 3 to 6 months, costing $100K - $500K.

Contact(s): Patricia E. Wirth (732.420.2100) will provide assistance from the initial client discussion through proposal preparation and project delivery.

Scalability Evaluation of Vendor Products

Description: Highly skilled engineers review vendor contracts and evaluate performance and scalability information provided by vendors. The team recommends, analyzes and performs necessary laboratory tests in order to ensure that a product performs and scales as expected with increasing load. They assist vendor managers in ensuring that performance and scalability criteria are part of the vendor contract, and participate in reviews with potential vendors to ensure that they have done due diligence that their product performs and scales in an efficient manner.

Profile of Prospective Customers: Typical buyers are senior development executives who want to ensure that performance and scalability are part of their purchasing decisions.

Value to Customer: The benefits are more informed decision making and the ability to avoid costly surprises and rework that result when a product does not perform or scale in accordance with the needs of the business.

Engagement Duration/Price: Typical projects take 2 to 3 months, costing $100K - $200K.

Contact(s): Patricia E. Wirth (732.420.2100) will provide assistance from the initial client discussion through proposal preparation and project

Network Scalability & Disaster Prevention

Description: A team of highly skilled and experienced engineers analyze the customer's Frame Relay, ATM, IP or Optical network (using an OSPF or PNNI type link-state protocol) in terms of its ability to survive a control message storm caused by the failure/recovery of trunks/nodes. They develop plans for the prevention of, and recovery from disasters based on proactive monitoring of critical network resources and taking corrective action at the onset of a message storm. Deliverables include detailed short and long-term recommendations on how to maintain network scalability and stability

Profile of Prospective Customers: Typical buyers are senior network engineering executives, CTOs and CIOs who want to ensure the stability of their networks even as they grow large.

Value to Customer: The benefit is in avoiding a major disaster that affects customers for an extended period of time (hours or days). The potential savings in lost revenue and customer goodwill is in the range of $10M to $100M per disaster event.

Engagement Duration/Price: Typical projects take 3 – 12 months, costing $300K - $900K.

Contact(s): Patricia E. Wirth (732.420.2100) will provide assistance from the initial client discussion through proposal preparation and

Traffic Engineering for Cable/DSL Access

Description: A cross-discipline team of highly skilled and experienced engineers develops and applies a methodology for optimally integrating voice, data and video services over Cable and/or DSL access networks. The design satisfies the Quality of Service needs of each user/application type and also becomes cost effective to the service provider. Deliverables are the Cable/DSL access network design and a final report containing detailed short and long-term recommendations.

Profile of Prospective Customers: Typical customers are senior network engineering executives, CTOs and CIOs who proactively seek to gain control of their Cable/DSL access investment.

Value to Customer: The benefit is a reduction in the amount of scarce upstream and downstream transmission bandwidth as well as the equipment needed at the customer premises and end-offices. The potential saving is in the range of $10M to $100M per year.

Engagement Duration/Price: Typical projects take 3 to 12 months, costing $300K - $900K.

Contact(s): Patricia E. Wirth (732.420.2100 will provide assistance from the initial client discussion through proposal preparation and project delivery.

Traffic Characterization & QoS Design

Description: A team of highly experienced telecommunications analysts and engineers will help the client define and set goals and objectives for introducing QoS to their network, identify what applications consume most of their network and server resources, and design an effective QoS solution that will deliver the desired level of performance to different application classes. The solution will recommend the QoS and traffic management controls as well as the monitoring and alerting tools. Deliverables include reports with detailed analysis and explicit recommendations.

Profile of Prospective Customers: The typical decision maker is the CIO organization or the telecommunication manager for the enterprise. Other customers would be network planning and network configuration divisions.

Value to Customer: The benefits are higher utilization and more economical use of the network, better performance management of the different applications, and more effective resource allocation among the organizations sharing the network.

Engagement Duration/Price: Typical projects take 3 to 12 months, costing $150K - $500K.

Contact(s): Patricia E. Wirth (732-420-2100) will provide assistance